

Shadowing Lemma and Chaotic Orbit Determination

Federica Spoto · Andrea Milani

5 June 2015

Abstract Orbit determination is possible for a chaotic orbit of a dynamical system, given a finite set of observations, provided the initial conditions are at the central time. In a simple discrete model, the standard map, we tackle the problem of chaotic orbit determination when observations extend beyond the predictability horizon. If the orbit is hyperbolic, a shadowing orbit is computed by the least squares orbit determination. We test both the convergence of the orbit determination iterative procedure and the behaviour of the uncertainties as a function of the maximum number n of map iterations observed. When the initial conditions belong to a chaotic orbit, the orbit determination is made impossible by numerical instability beyond a computability horizon, which can be approximately predicted by a simple formula. Moreover, the uncertainty of the results is sharply increased if a dynamical parameter is added to the initial conditions as parameter to be estimated. The uncertainty of the dynamical parameter decreases like n^a with $a < 0$ but not large (of the order of unity). If only the initial conditions are estimated, their uncertainty decreases exponentially with n . If they belong to a non-chaotic orbit the computational horizon is much larger, if it exists at all, and the decrease of the uncertainty is polynomial in all parameters, like n^a with $a \simeq 1/2$. The Shadowing Lemma does not dictate what the asymptotic behaviour of the uncertainties should be. These phenomena have significant implications, which remain to be studied, in practical problems of orbit determination involving chaos, such as the chaotic rotation state of a celestial body and a chaotic orbit of a planet-crossing asteroid undergoing many close approaches.

Keywords Chaotic Motions · Numerical Methods

1 Introduction

Chaotic dynamical systems are characterized by the existence of a *predictability horizon* in time, beyond which the information on the state available from the initial conditions is not enough for meaningful predictions. Thus it appears a difficult task to perform

F. Spoto · A. Milani
Dipartimento di Matematica, Università di Pisa, Largo Pontecorvo 5, 56127 Pisa, Italy E-mail: spoto@spacedys.com

an orbit determination for a chaotic dynamical system, at least when the observations are spread over a time-span longer than the predictability horizon.

Nevertheless there are practical problems of orbit determination in which the system is chaotic and the time-span of the observations is very long. It is important to understand the behaviour of the solutions, with their estimated uncertainties, in particular when the variables to be solved for include not just the initial conditions but also some dynamical parameters. If the number of available observations grows, but simultaneously the time interval over which they are spread grows up to values comparable to the predictability horizon, does the solution become more accurate, and is the iterative procedure of differential corrections (Milani and Gronchi, 2010, Chap. 5) to find the least squares solution still possible?

In this paper we use a model problem, namely the discrete dynamical system defined by the standard map of the pendulum, with just one dynamical parameter, the μ coefficient appearing in equation (1). We also set up an observation process in which both coordinates of the standard map are observed after each iteration. In the observations we include a simulated random noise with a normal distribution. Then, each experiment of orbit determination is also a concrete computation of a segment limited to n iterations (of the map and of its inverse) of a ε *shadowing orbit* for the δ *pseudo trajectory* defined by the observation process. The *Shadowing Lemma* (see Section 3) provides a mathematically rigorous result on the availability of shadowing orbits, but thanks to the orbit determination process we make explicit the relationship between ε and δ (see Section 3.1), and we explicitly compute the ε -shadowing orbit.

At the same time, each experiment provides an estimate of the standard deviation of each of the variables, including initial conditions and the parameter. These estimates can be plotted as a function of n , thus showing the relationship between accuracy, number of observations and time interval, measured in Lyapounov times (see Section 4).

Of course the numerical experiments are limited to a finite number of iterations, while the Shadowing Lemma refers to an infinite orbit. However, the maximum number of iterations is controlled by another time limit, the *computability horizon* due to round off error. This limit can be estimated approximately by a simple formula, and it is found in numerical experiments as a function of both the initial conditions and the numeric precision used in the computations.

1.1 Wisdom hypothesis

In 1987 J. Wisdom was discussing the chaotic rotation state of Hyperion, when he claimed that numerical experiments indicated that *the knowledge gained from measurements on a chaotic dynamical system grows exponentially with the timespan covered by the observations* (Wisdom, 1987). This pertained in particular to the information on dynamical parameters like the moments of inertia ratios for Hyperion, as well as the rotation state at the midpoint of the time interval covered by the observations, which he proposed would be determined with exponentially decreasing uncertainty.

Therefore Wisdom suggests that the orbit determination for a chaotic system might be in fact more effective than for a non-chaotic one. It is clear from the context that he was referring to numerical results, thus his statement can only be verified with finite computations as close as possible to a realistic data processing of observations of a chaotic system with dynamical parameters to be determined.

We have set as a goal in this paper to test the behavior of the uncertainty in the dynamical parameter of our model problem. We shall discuss the implications for Wisdom's claim in Section 5.1.

1.2 Application to planet-crossing asteroids

In our solar system there are *planet-crossing minor bodies*, including asteroids and comets, by definition such that their orbits can, at some times, intersect the orbit of the major planets. In particular many of the *Near Earth Asteroids (NEAs)* can intersect the orbit of the Earth. These orbits are necessarily chaotic, at least over the timespan accessible to accurate numerical computations.

Unfortunately, these orbits are especially important and necessary to be studied because of the very reason of chaos, namely close approaches to the major planets including the Earth: these approaches may, in some cases, be actual impacts on a finite size planet.

The attempt to predict possibility of impacts by NEAs, in particular on our planet, is called *Impact Monitoring*, and it is indeed a form of orbit determination for chaotic orbits. There is a subset of cases of NEAs for which non-gravitational perturbations, such as the ones resulting from the Yarkovsky effect, are not negligible in terms of Impact Monitoring because of the exponential divergence of nearby orbits which amplifies initially very small perturbations

(Farnocchia et al., 2013; Chesley et al., 2014; Farnocchia and Chesley, 2014; Spoto et al., 2014).

Thus the Impact Monitoring for these especially difficult cases is an instance of orbit determination of a chaotic system, with as parameters the 6 initial conditions and at least one dynamical parameter, such as a Yarkovsky effect coefficient to be solved for. We shall show in Section 5.2 that the weak determination of the dynamical parameter is a key feature of these cases.

2 Orbit determination for the standard map

The simplest example of a conservative dynamical system which has both chaotic and ordered orbits can be built by means of an area preserving map of a two dimensional manifold:

$$S_\mu(x_k, y_k) = \begin{cases} x_{k+1} = x_k + y_{k+1} \\ y_{k+1} = y_k - \mu \sin x_k \end{cases} \quad (1)$$

where μ is the perturbation parameter, and S is the standard map. The system has more regular orbits for small μ , and more chaotic orbits for large μ . We choose an intermediate value of μ , e.g. $\mu = 0.5$, in such a way that ordered and chaotic orbits are both present. Figure 1 shows the strongly chaotic region around the hyperbolic fixed point, and a few regular orbits on both sides.

The advantage of such example is that the least square parameter estimation process can be performed by means of an explicit formula.

First we compute the linearized map

$$DS = \begin{pmatrix} \frac{\partial x_{k+1}}{\partial x_k} & \frac{\partial x_{k+1}}{\partial y_k} \\ \frac{\partial y_{k+1}}{\partial x_k} & \frac{\partial y_{k+1}}{\partial y_k} \end{pmatrix} = \begin{pmatrix} 1 - \mu \cos(x_k) & 1 \\ -\mu \cos(x_k) & 1 \end{pmatrix}$$




Fig. 1 Orbits of the standard map for the perturbation parameter $\mu = 0.5$. Plotted is a blow up of the central region around the hyperbolic fixed point, showing the strongly chaotic region and a few regular orbits on both sides.

and from this the linearized state transition matrix

$$A_k = \frac{\partial(x_k, y_k)}{\partial(x_0, y_0)}$$

which is the solution of the variational equation (for infinitesimal displacement in the initial conditions), and given by the recursion:

$$A_{k+1} = DS A_k \quad ; \quad A_0 = I .$$

The variational equation for the derivatives with respect to the model parameter μ is:

$$\begin{aligned} \frac{\partial(x_{k+1}, y_{k+1})}{\partial\mu} &= DS \frac{\partial(x_k, y_k)}{\partial\mu} + \frac{\partial S_\mu}{\partial\mu} \\ &= DS \frac{\partial(x_k, y_k)}{\partial\mu} + \begin{pmatrix} -\sin(x_k) \\ -\sin(x_k) \end{pmatrix} \end{aligned}$$

Then we set up an observation process, in which both coordinates x and y are observed at each iteration, and the observations are Gaussian random variables with mean x_k (y_k respectively) and standard deviation σ : we use the notation $x_k(\mu_0, \sigma)$ to indicate that the probability density function of the observation x_k is the normal $\mathcal{N}(x_k(\mu_0), \sigma^2)$ one, and similarly for y_k . The residuals are:

$$\begin{cases} \xi_k = x_k(\mu_0, \sigma) - x_k(\mu_1) \\ \bar{\xi}_k = y_k(\mu_0, \sigma) - y_k(\mu_1). \end{cases} \quad (2)$$

for $k = -n, \dots, n$. In (2) $x_k(\mu_0, \sigma)$ and $y_k(\mu_0, \sigma)$ are the observations at each iteration, μ_0 is the “true” value and $\mu_1 = \mu_0 + d\mu$ is the current guess.

Then the least squares fit is obtained from the normal equations (Milani and Gronchi, 2010):

$$C = \sum_{k=-n}^n B_k^T B_k ; \quad D = - \sum_{k=-n}^n B_k^T \begin{pmatrix} \xi_k \\ \bar{\xi}_k \end{pmatrix} \quad (3)$$

$$B_k = \frac{\partial(\xi_k, \bar{\xi}_k)}{\partial(x_0, y_0, \mu)} = - \left(A_k \middle| \frac{\partial(x_k, y_k)}{\partial\mu} \right).$$

The least squares solution for both, the parameter μ and the initial conditions, is:

$$\begin{pmatrix} \delta x \\ \delta y \\ \delta \mu \end{pmatrix} = \Gamma D, \quad \Gamma = C^{-1}$$

with Γ the covariance matrix of the result. This is enough to find the least squares solution by iteration of the above *differential correction*. However, to assess the uncertainty of the result, weights should be assigned to the residuals consistently with the probabilistic model, in this case each residual needs to be divided by its standard deviation σ ; then the distribution of the result (x, y, μ) is a normal distribution with covariance matrix $\sigma^2 \Gamma$.

3 Shadowing Lemma

The shadowing problem is that of finding a deterministic orbit as close as possible to a given noisy orbit. The so-called Shadowing Lemma is the main result about shadowing near a hyperbolic set of a diffeomorphism. Anosov (Anosov, 1967) and Bowen (Bowen, 1975) proved the existence of arbitrarily long shadowing solutions for invertible hyperbolic maps. Here we give an overview of these classical results, as in (Pilyugin, 1999).

Let (X, d) be a metric space and let Φ be a homeomorphism mapping X onto itself. A δ -pseudotrajectory of the dynamical system Φ is a sequence of points $\zeta = \{x_k \in X : k \in \mathbb{Z}\}$ such that the following inequalities

$$d(\Phi(x_k), x_{k+1}) < \delta. \quad (4)$$

hold. For a graphical description of a δ -pseudotrajectory, see Fig. 2. Usually, a δ -




Fig. 2 A δ -pseudotrajectory.

pseudotrajectory is considered as the result of using a numerical method to compute




Fig. 3 An ε -shadowing.

orbits of the dynamical system Φ , e.g., because of round off error. We say that a point $x \in X$ (ε, Φ) -*shadows* a pseudotrajectory $\zeta = \{x_k\}$ if the inequalities

$$d(\Phi^k(x), x_k) < \varepsilon \quad (5)$$

hold (see Figure 3). If only one dynamical system Φ is considered, we will usually write ε -shadows ζ . The existence of a shadowing point for a pseudotrajectory ζ means that ζ is close to a real trajectory of Φ .

The following statement is usually called the Shadowing Lemma.

Theorem 1 *If Λ is a hyperbolic set for a diffeomorphism Φ , then there exists a neighborhood W of Λ such that for all $\varepsilon > 0$ there exists $\delta > 0$ such that for any δ -pseudotrajectory with initial conditions $\zeta \in W$ there is a point x that ε -shadows ζ .*

The Anosov shadowing requires the existence of a hyperbolic set. It means that at each point there are two directions where the motion is either exponentially expanding (stable manifold) or exponentially contracting (unstable manifold).

Definition 1 We say that a set Λ is hyperbolic for a diffeomorphism $\Phi \in C^1(\mathbf{R}^n)$ if:

- (a) Λ is compact and Φ -invariant;
- (b) there exist constants $C > 0$, $\lambda_0 \in (0, 1)$, and families of linear subspaces $S(p)$, $U(p)$ of \mathbf{R}^n , $p \in \Lambda$, such that
 - (b.1) $S(p) \oplus U(p) = \mathbf{R}^n$;
 - (b.2) $D\Phi(p)T(p) = T(\Phi(p))$, $p \in \Lambda$, $T = S, U$;
 - (b.3)

$$\begin{aligned} |D\Phi^m(p)v| &\leq C\lambda_0^m|v| \text{ for } v \in S(p), \ m \geq 0; \\ |D\Phi^{-m}(p)v| &\leq C\lambda_0^m|v| \text{ for } v \in U(p), \ m \geq 0; \end{aligned}$$

The definition of a hyperbolic set is strictly related to the one of Lyapounov exponent: for each orbit inside a hyperbolic set, the Lyapounov exponents must be either $> \log(\lambda_0)$ or $< -\log(\lambda_0)$.

3.1 Shadowing Lemma and orbit determination

Our goal is to connect the Shadowing Lemma with the chaotic orbit determination, involving the least squares fit and the differential corrections.

First of all we build a δ -pseudotrajectory by using a simulated observations process. In Section 2 we have set up such an observations process, with observations

$(x_k(\mu_0, \sigma), y_k(\mu_0, \sigma))$. We claim that a sequence $\zeta = \{(x_k(\mu_0, \sigma), y_k(\mu_0, \sigma))\}$ is a δ -pseudotrajectory for the dynamical system $S_{\mu^*}(x_0, y_0)$, with $\delta = \sqrt{2}|\mu^* - \mu_0| + \mathcal{K}\sigma$, $\mathcal{K} \in \mathbf{R}$. To obtain this result we compute the Euclidean distance:

$$d(S_{\mu^*}(x_k(\mu_0, \sigma), y_k(\mu_0, \sigma)), (x_{k+1}(\mu_0, \sigma), y_{k+1}(\mu_0, \sigma))) \quad (6)$$

For the sake of simplicity $(\bar{x}_{k+1}, \bar{y}_{k+1})$ are the observations, i.e. Gaussian random variables with mean x_{k+1} (y_{k+1} , respectively), and standard deviation σ , as in Sec. 2, and $S_{\mu^*}(\bar{x}_k, \bar{y}_k) = (\bar{x}_{k+1}, \bar{y}_{k+1})$. Using these notations, (6) turns into

$$d(S_{\mu^*}(\bar{x}_k, \bar{y}_k), (\bar{x}_{k+1}, \bar{y}_{k+1})) = \sqrt{(\bar{x}_{k+1} - \bar{x}_{k+1})^2 + (\bar{y}_{k+1} - \bar{y}_{k+1})^2}$$

We compute separately the two differences.

$$\begin{aligned} |\bar{y}_{k+1} - \bar{y}_{k+1}| &= |\bar{y}_{k+1} - \mu^* \sin \bar{x}_k - y_{k+1} - \mathcal{N}(0, \sigma^2)| \\ &= |\mathcal{N}(0, 2\sigma^2) - \mu^* \sin x_k \cos(\mathcal{N}(0, \sigma^2)) + \mu^* \sin(\mathcal{N}(0, \sigma^2)) \cos x_k + \mu_0 \sin x_k| \\ &< \mathcal{N}(0, 2\sigma^2) + |\mu_0 - \mu^*| \end{aligned} \quad (7)$$

To allow the last estimate, we need to solve a technical problem: the Shadowing Lemma uses a uniform norm, that is the maximum of the distance between the δ -pseudotrajectory and the ε -shadowing. On the contrary, the natural norm for the residuals of the fit is the Euclidean norm with the square root of the sum of the squares. However, since the number of residuals is not only finite but sharply limited by the numerical phenomena discussed in Section 4, in a given experiment we can just take the maximum absolute value of the residuals and it is going to be $\mathcal{K}\sigma$, with \mathcal{K} a number which in practice cannot be too large, e.g., in our experiment $\mathcal{K} = 5.9$.

Then we can approximate the quantities $\mathcal{O}(\sigma)$ and smaller, e.g., $\cos(\mathcal{N}(0, \sigma^2)) \sim 1$ and $\sin(\mathcal{N}(0, \sigma^2)) \sim 0$. The x coordinate gives a similar result:

$$\begin{aligned} |\bar{x}_{k+1} - \bar{x}_{k+1}| &= |\bar{x}_{k+1} + \bar{y}_{k+1} - x_k - y_{k+1} - \mathcal{N}(0, \sigma^2)| \\ &< \mathcal{N}(0, \sigma^2) + \mathcal{N}(0, 2\sigma^2) + |\mu_0 - \mu^*| = \mathcal{N}(0, 3\sigma^2) + |\mu_0 - \mu^*| \end{aligned} \quad (8)$$

Putting together (7) and (8) we obtain

$$\begin{aligned} &\sqrt{(\bar{x}_{k+1} - \bar{x}_{k+1})^2 + (\bar{y}_{k+1} - \bar{y}_{k+1})^2} < \\ &< \sqrt{(\mathcal{N}(0, 3\sigma^2) + |\mu_0 - \mu^*|)^2 + (\mathcal{N}(0, 2\sigma^2) + |\mu_0 - \mu^*|)^2} < \quad (9) \\ &< \sqrt{2}|\mu_0 - \mu^*| + \sqrt{(\mathcal{N}(0, 3\sigma^2))^2 + (\mathcal{N}(0, 2\sigma^2))^2} < \sqrt{2}|\mu_0 - \mu^*| + \mathcal{K}\sigma \end{aligned}$$

with $\mathcal{K} \in \mathbf{R}$.

Therefore the sequence generated by the observations is a δ -pseudotrajectory for the dynamical system S_{μ^*} with $\delta = \sqrt{2}|\mu_0 - \mu^*| + \mathcal{K}\sigma$.

Figure 4 is an example of observations as a δ -pseudotrajectory. The observations are built with $\sigma = 10^{-3}$ and $\mu_0 = 0.5$, and the dynamical system is S_{μ^*} , with $|\mu^* - \mu| = 10^{-1}$; the circles have radius δ .

The solution of the least squares fit (to the observations from $-n$ to n), obtained by convergent differential corrections, is a finite ε -shadowing, valid for the iterations from $-n$ to n .

We choose a value $\varepsilon > 0$, that is a boundary on the maximum of the residuals. Then we choose $\sigma < \varepsilon/\mathcal{K}$ and we set up the observations process. Next, we create




Fig. 4 An example of a δ -pseudotrajectory. Initial conditions are $x_0 = 3$, $y_0 = 0$, $\mu_0 = 0.5$. $\delta\mu = 10^{-1}$, and $\sigma = 10^{-3}$.

a first guess: a new orbit obtained with a small change of the initial conditions and of the dynamical parameter μ : $\{(x_k(\mu_g), y_k(\mu_g))\} = S_{\mu_g}(x_g, y_g)$, with $x_g = x_0 + dx$, $y_g = y_0 + dy$, and $\mu_g = \mu_0 + d\mu$. Then we apply the differential corrections to the orbit. If the iterations converge, that is the last correction is very small, the maximum of the norm of the residuals is less than ε (because the individual residuals are less than 3σ).

At convergence, we obtain an initial condition (x^*, y^*) and a value of the dynamical parameter μ^* , such that (x^*, y^*) is the (ε, S_{μ^*}) -shadowing for the δ -pseudotrajectory with $\delta = \sqrt{2}|\mu^* - \mu_0| + \mathcal{K}\sigma$, for all the points used in the fit.

The most important requirement is the convergence of the differential corrections, otherwise we cannot obtain the ε -shadowing. This is far from trivial, because the chaotic divergence of the orbits introduces enormous nonlinear effects, for which the linearized approach of differential corrections may fail. To guarantee convergence, first we select the initial conditions x_0, y_0 to be at the center of the observed interval, otherwise the initial conditions would be essentially undetermined along the stable manifold of the initial conditions. Second, we use a *progressive solution* approach, namely, given the solution with $2n+1$ observations with indexes between $-n$ and n , we use the convergent solution x_0^*, y_0^*, μ^* for n as first guess for the solution with $2n+3$ observations (between $-n-1$ and $n+1$). Then the initial guess is actually used only for the solution with 3 observations, for which the nonlinearity is negligible. Still the convergence of the differential corrections depends critically upon the number n of iterations of the map, as explained in Section 4.

4 Numerical results

The experiment was performed with the initial conditions at $x_0 = 3$ and $y_0 = 0$, and the value of the dynamical parameter $\mu_0 = 0.5$. The dynamical context for this orbit can be appreciated from Figure 1, showing that the initial conditions are indeed in a portion of the initial conditions space containing mostly chaotic orbits. For the observation noise we have used a standard deviation $\sigma = 10^{-10}$.

4.1 Computability horizon




Fig. 5 The eigenvalues and the determinant of the state transition matrix in a semilogarithmic scale, as a function of the number of iterations. Also shown is the linear fit to the large eigenvalue based on the first 180 iteration, with slope $+0.091$. The computation is in double precision and the number of iterations of the standard map n is 300 with the map and 300 with its inverse. The determinant of the state transition matrix would be 1, for all n , in an exact computation. The numerical instability occurs when the eigenvalues reach the critical values $\sqrt{\varepsilon_d}$, $\sqrt{1/\varepsilon_d}$ marked by the dotted lines.

Figures 5 and 6 show the absolute value of the eigenvalues of the state transition matrix forward and backward. The product of two eigenvalues should be 1 in exact arithmetic. When the condition number of the matrix becomes larger than the inverse




Fig. 6 The eigenvalues and the determinant of the state transition matrix in a semilogarithmic scale, as a function of the number of iterations. Also shown is the linear fit to the large eigenvalue based on the first 300 iteration, with slope $+0.086$. The computation is in quadruple precision and the number of iterations of the standard map n is 800 with the map and 800 with its *inverse*. The numerical instability occurs when the eigenvalues reach the critical values $\sqrt{\varepsilon_q}$, $\sqrt{1/\varepsilon_q}$ marked by the dotted lines.

of the machine rounding off error, the computation of the matrix becomes numerically impossible, and the computed value of the determinant is far from 1.

In Figure 5 the computations are performed in the standard double precision, that is with a mantissa of 52 binary digits and a round off relative error of $\varepsilon_d = 2^{-53} = 1.1 \times 10^{-16}$. We observe a numerical instability after $\simeq 180$ iterations: the determinant deviates from the exact value of 1 and the small eigenvalue starts increasing; the large eigenvalue keeps increasing, but there is a slight change of slope. Then we fit the slope of the large eigenvalue curve for the first 180 iterations, and get a Lyapounov indicator $+0.091$: it approximates the maximum Lyapounov exponent χ for the orbit to which our differential corrections converge¹.

¹ There is no way to rigorously compute the Lyapounov exponents: in practice *Lyapounov indicators* extracted from finite propagations are used to assess, but not rigorously prove, the chaotic nature of the orbits. Note that it is a numerically well documented phenomenon that the indicators are not constant, but actually depend upon the time interval over which they are computed, although in most cases these changes are not very large and the conclusion that an orbit is chaotic is reliable.

The Lyapounov time is $T_L = 1/\chi$, in this example $T_L \simeq 11$. To reach a ratio of eigenvalues of the state transition matrix of $1/\varepsilon_d$ we need a number of Lyapounov times $\log(1/\sqrt{\varepsilon_d})$, in this case $\simeq 18.4 T_L \simeq 202$ iterations of the map. At about this number of iterations the maximum and minimum eigenvalues of A_n are so widely apart in size that a bad conditioning horizon is reached, and the computation of the state transition matrix becomes numerically inaccurate. Hence near $\pm 18.4 T_L$ we observe the numerical instability in the computation of the determinant and of the eigenvalues of the state transition matrix.

Figure 6 shows the same computations, with the same initial condition, but in quadruple precision, with a 112 bit mantissa and $\varepsilon_q = 2^{-113} = 9.6 \times 10^{-35}$. The change of slope in the eigenvalues curves occurs after $\simeq 300$ iterations, while a full blown numerical instability occurs after $\simeq 550$ iterations. The fit to the large eigenvalue for the first 300 iterations gives a Lyapounov indicator $+0.086$, not very different from the one obtained in double precision. Thus we would expect the numerical instability to occur after $\log(1/\sqrt{\varepsilon_q}) T_L \simeq 39.2 T_L = 455$ iterations. It appears that the rate of divergence decreases after 300 iterations, as shown by the change in slope, allowing to maintain at least the determinant near the exact value for about 100 more iterations beyond the value predicted above.

The *computability horizon* represents the maximum number of iterations we can reach, before the computation becomes numerically unstable. The computability horizon strongly depends on the chaoticity of the system: more chaos, that is larger χ , more instability; but also upon the precision of the computations.

Thus, in the following we perform the numerical experiments in quadruple precision, in order to mitigate the problem of the numerical instability. We compute 500 iterations forwards and backwards, but we use only the first 300 iterations for the linear fits, to avoid the possibility that changes in slope, such as the ones apparent in Figure 6, contaminate our experimental results.

The computability horizon is a hard limit in that it is not practically possible to increase the number of iterations to a much higher value. E.g., to push the horizon by a factor 10 above the value for double precision, we would need computations performed with real numbers represented with 800 bytes².

The conclusion is that the practical problem of chaotic orbit determination is meaningful only for a finite number of iterations, and the accuracy of the results can be tested only within the boundary of the computability horizon.

4.2 Chaotic case

Figure 7 shows the results in quadruple precision of the full 3-parameter fit: the 3 parameters are the initial conditions and the dynamical parameter μ . The determination of μ is indeed not possible without simultaneous determination of the initial conditions.

Even in quadruple precision we find a maximum value of n beyond which the iterative solution of the nonlinear least squares problem is divergent. This maximum turns out to be 599 in this experiment: it is close to what we have called the computability horizon, that is this limitation is due to the difficulty of computing the state transition matrix when the condition number is too large.

² Software to perform arithmetic computations with an arbitrary number of digits is available, but the algorithms are too slow to be used even for our simple example.




Fig. 7 Standard deviation of the solutions for the initial conditions and for the dynamical parameter μ (continuous lines), and actual error (nominal solution of the fit minus real value used in the simulation, dashed lines), as a function of the number of iterations.

The curves in Fig. 7 represent both the formal standard deviation and the actual error of the solutions of the least squares fit, as a function of n in a semilogarithmic plot. Both the formal standard deviation and the actual error of μ do not decrease exponentially. Indeed, in Fig. 8 we have the same behavior of the curves that we have already seen in Fig. 7, but in a log-log plot, in which a constant slope a would imply a power law proportional to n^a . The slopes of the lines that fit the uncertainties are: -0.675 for the dynamical parameter μ , -0.833 and -12.030 for the initial conditions x and y , respectively. This plot in logarithmic scale shows that the uncertainty for μ and x does not decrease exponentially. It is also apparent that one of the initial conditions (y) is better determined than the other one (x), with an improvement as a function of n which could be exponential. This is a property of the specific initial condition we have used, for other choices we can get three parameters determined with comparable accuracy, none of them with exponential improvement³.

Figure 9 shows the results for the standard deviation and the actual error when solving only for the initial condition. The 2×2 portion of the normal matrix which refers only to the initial conditions is not badly conditioned. Also as a result of this, we are able to get convergence of the differential corrections up to ± 742 iterations, which is

³ This depends upon the orientation of the stable and unstable directions at the initial condition.




Fig. 8 Uncertainty of the solution of the least squares fit for the initial conditions and for the dynamical parameter μ in a logarithmic scale.

even beyond the numerical stability boundary. If the fit is done by using only up to 300 iterates, to avoid the apparent slope change, the slopes shown in this Figure are -0.084 for x and -0.083 for y ; note that the Lyapounov indicator for the same interval is $+0.086$. Thus exponentially improving determination of the initial conditions only is possible, and the exponent appears to be very close to the opposite of the Lyapounov exponent.

4.3 Ordered case

An ordered case can be obtained with a change of the initial conditions. For the numerical experiments we have chosen $x_0 = 2$, $y_0 = 0$ and $\mu_0 = 0.5$. In the ordered case we have not the problem of the computability horizon and the Lyapounov exponent is very small: actually, it could be zero if we are on a Moser invariant curve. Thus we have computed 5000 iterations.

Figure 10 gives a summary of our numerical experiment in the ordered case. The Lyapounov indicator is very small ($\simeq 10^{-4}$), and can be made even smaller by continuing the experiment for larger values of n . As a consequence, the state transition matrix is not badly conditioned, and the computability horizon is much beyond the number of iterations we have used (if it exists at all). Thus the lack of chaoticity implies the




Fig. 9 Standard deviation of the solutions for the initial conditions (continuous lines), true errors for the same 2 parameters (dashed lines).

practical absence of the computability horizon, and we can determine all the parameters with very good accuracy, even if we are not in exact arithmetic. The values of the slopes of the fit to the uncertainty are -0.504 for μ , -0.504 and -0.488 for x and y respectively, the corresponding regression lines are shown in the log-log plot on the bottom right. As it is clear by comparing the top right and the bottom left plot, the standard deviation for the solution with only 2 parameters have very much the same behavior, indeed in a log-log plot (not shown) we can get slopes -0.511 for x , -0.481 for y .

All these power laws are close to the inverse square root of the number of iterations, namely the same rule as the standard deviation in the computation of a mean. We do not have a formal proof of this, but we conjecture that for an orbit on a Moser invariant curve (for which the Lyapounov exponents are exactly zero) the standard deviations for all the parameters decrease as $1/\sqrt{n}$.

5 Conclusions

We have understood the concept of *computability horizon* as a consequence of numerical instability in the computation of the state transition matrices, thus providing a comparatively simple empirical formula to approximately predict the horizon. This is

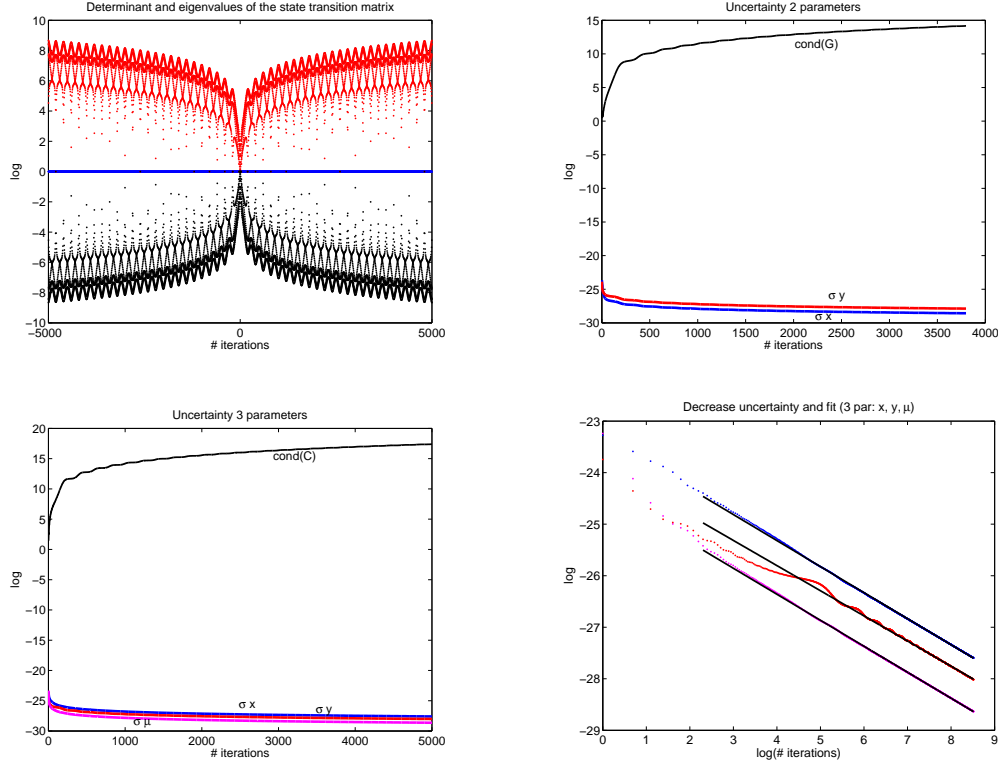


Fig. 10 Top left: eigenvalues of the state transition matrices, for the chosen regular initial conditions and for ± 5000 iterations. Top right: solutions for the initial condition only, from the top condition number of the normal matrix, standard deviation of y and for x . Bottom left: solutions for three parameters, from the top condition number, standard deviation of x , of y , of μ . Bottom right: log-log plot of the 3 standard deviations, with very similar slopes.

a practical limitation which applies to any attempt at orbit determination of chaotic orbits.

We have used numerical experiments in quadruple precision, but nevertheless limited by the computability horizon to few hundreds of iterations for the very chaotic orbit used in our test. From these we have found the following three empirical facts:

1. If only initial conditions are determined for a chaotic orbit, the uncertainty can decrease exponentially with the number n of iterations of the map, and the exponent of this decrease is close to a Lyapounov exponent.
2. If a dynamical parameter μ is determined together with the initial conditions of a chaotic orbit, the decrease in uncertainty is polynomial in n for μ and for at least one of the initial coordinates.
3. If the initial conditions belong to an apparently ordered orbit, that is such that there is no evidence of a positive Lyapounov exponent, it is possible to determine simultaneously μ and the initial conditions with uncertainty decreasing polynomially with n . The case in which only the initial conditions are determined gives the same

result. Moreover, all the power laws n^a for these uncertainties appear numerically to have $a \simeq -1/2$.

5.1 Discussion on the Wisdom hypothesis

The statement by Wisdom, as a practical rule for concrete orbit determination, appears to be first limited by the computability horizon. Second, the actual decrease of the uncertainty, going as far as it can be done numerically, is not exponential, but polynomial, as n^a , with a negative and rather small, although we have found that the value of a depends upon the initial conditions⁴. Note that the orbit determinations in which the only parameters to be solved are the 2 initial coordinates show an exponential decrease as $\exp(-\alpha n)$, where α appears to be close to the Lyapounov exponent χ , but the strong correlations appearing when 3 parameters are solved degrade the result in a very substantial way.

This needs to be compared to the regular case, shown in Figure 10, where the standard deviations for each of the 3 fit parameters decrease approximately according to an $1/\sqrt{n}$ law, as prescribed by the standard rule for the estimate of the mean with errors having a normal distribution. Indeed it is possible that the determination of μ for some chaotic cases, including the example shown in Figure 8, decreases faster than for an ordered case, but the decrease is anyway polynomial, proportional to n^a with some different negative a , thus the difference is not very large, given the tight constraint on the maximum possible value of n .

5.2 Examples from Impact Monitoring

One feature of our results is that adding a dynamical parameter to the list of parameters to be determined results in degradation in the normal matrix, thus in much slower decrease of the uncertainties as the number of observations grows. The problems of orbit determination for NEA undergoing several close approaches to the Earth (or other planets) is more complex than our simple model, but we have found that the phenomenon described above does occur in a remarkably similar way.

In Figure 11 we show two probability distributions, as derived from the orbit determination of the asteroid (410777) 2009 FD. The narrow peaked distribution corresponds to an orbit determination with 6 parameters, the initial conditions only: the standard deviation is 6×10^4 km. The much wider distribution corresponds to a fit with 7 parameters, including the constant A_2 appearing in the transverse acceleration due to the Yarkovsky effect: the STD is $\simeq 2.3 \times 10^6$ km. The Yarkovsky effect is a form of non-gravitational perturbation due to thermal radiation emitted anisotropically by the asteroid, and is indeed very small. However, when the uncertainty resulting from the covariance matrix of the orbit determination is propagated for ~ 170 years after the last observation available, not only the Yarkovsky effect has a long enough time to accumulate but it is also enhanced by the exponential divergence of nearby orbits, the Lyapounov time being about 15.3 years (Spoto et al., 2014)[Figure 5].

⁴ We are showing figures and giving data only for one initial condition, but of course we have run many tests.




Fig. 11 Two different Probability Density Functions (PDF) for the trace of possible solutions on the Target Plane of the close approach of asteroid (410777) 2009 FD to the Earth in the year 2185. Superimposed and on a different vertical scale are the keyholes relative to impacts in different years between 2185 and 2196; the height of the bar is proportional to the width of the keyhole, thus the Impact Probability can be computed as product of the width and the PDF.

The practical consequence of this increase of the uncertainty arises from the fact that the Target Plane of 2009 FD for 2815 includes some *keyholes*, small portions corresponding to impacts with the Earth (either at that time or a few years later, until 2196). With the 7 parameters solutions these keyholes are within the range of outcomes with a significant value of the Probability Density Function, thus the impacts have a non-negligible probability, the largest being an *Impact Probability* of $\simeq 1/370$ for 2185. If on the contrary the orbit was estimated with 6 parameters only, then the probability would appear to be even larger for an impact in 2190, and all the other keyholes (including the one for 2185) would correspond to negligible impact probabilities. Given that the impact, if it was to occur, would release an energy equivalent to 3,700 Mega-Tons of TNT, this difference is practically relevant. In fact, the solution including the Yarkovsky effect leads to a more reliable estimate of the Impact Probabilities, because the Yarkovsky effect exists and needs to be taken into account.

Is the discrepancy in the uncertainties with and without the dynamical parameter in the fit essentially the same phenomenon we have found in our simple model? We do not know the answer to this question, but we shall investigate this issue in the future.

Acknowledgements This work has been partially supported by the Marie Curie Initial Training Network Stardust, FP7-PEOPLE-2012-ITN, Grant Agreement 317185. The authors were also sponsored by an internal research fund of the Department of Mathematics of the University of Pisa.

References

- Anosov, D., 1967. Geodesic flows on closed Riemann manifolds with negative curvature. In: Proc. Steklov Inst. Math, 90.
- Bowen, R., 1975. ω -limit sets for axiom A diffeomorphism. J. Diff. Eqs. 18, 333–356.
- Chesley, S. R., Farnocchia, D., Nolan, M. C., Vokrouhlický, D., Chodas, P. W., Milani, A., Spoto, F., Rozitis, B., Benner, L. A. M., Bottke, W. F., Busch, M. W., Emery, J. P., Howell, E. S., Lauretta, D. S., Margot, J.-L., Taylor, P. A., Jun. 2014. Orbit and bulk density of the OSIRIS-REx target Asteroid (101955) Bennu. Icarus 235, 5–22.
- Farnocchia, D., Chesley, S. R., Feb. 2014. Assessment of the 2880 impact threat from Asteroid (29075) 1950 DA. Icarus 229, 321–327.
- Farnocchia, D., Chesley, S. R., Chodas, P. W., Micheli, M., Tholen, D. J., Milani, A., Elliott, G. T., Bernardi, F., May 2013. Yarkovsky-driven impact risk analysis for asteroid (99942) Apophis. Icarus 224, 192–200.
- Milani, A., Gronchi, G. F., 2010. Theory of Orbital Determination. Cambridge University Press.
- Pilyugin, S., 1999. Shadowing in Dynamical Systems. Springer-Verlag.
- Spoto, F., Milani, A., Farnocchia, D., Chesley, S. R., Micheli, M., Valsecchi, G. B., Perna, D., Hainaut, O., Dec. 2014. Nongravitational perturbations and virtual impactors: the case of asteroid (410777) 2009 FD. Astronomy & Astrophysics 572, A100.
- Wisdom, J., Nov. 1987. Urey Prize Lecture - Chaotic dynamics in the solar system. Icarus 72, 241–275.